

Thursday, June 21, 2001

Birth of a Thinking Machine

■ For 17 years, a team has been trying to develop the most sophisticated artificial intelligence system ever. This summer, the public will be able to see its work.

By MICHAEL A. HILTZIK, Times Staff Writer

AUSTIN, Texas--Popular culture has long held strong opinions about what the world's smartest machine should look like. There's the unblinking red eye of HAL, the brilliant, homicidal computer of Stanley Kubrick's "2001: A Space Odyssey"; the gilded humanoids of pulp sci-fi; and the flashing lights and gleaming boxes of countless doomsday scenarios.

But it's a safe bet that nobody has imagined artificial intelligence the way it is taking shape inside a low-slung brown brick building hidden deep within a leafy research park north of town. Yet here beats the heart of the system known as Cyc.

For 17 years a small band of engineers and programmers has been slaving away at the task of teaching Cyc much of what a human being knows. (The name comes from "encyclopedia" and is pronounced "psych.") The idea, as articulated by the project's creator, computer scientist Douglas B. Lenat, has been to create the most sophisticated artificial intelligence system ever devised--the closest a computer has come to replicating the human brain's reasoning, learning ability and perhaps even its consciousness.

Whether Lenat and his team have achieved that goal is about to be subjected to public scrutiny. For years, Cyc has largely been kept under wraps, seen and used only by specialists and a handful of commercial firms and government agencies licensed by Cyc's corporate owner, Cycorp Inc.

But Cyc is about to have its coming-out party. Late this summer (soon after Steven Spielberg's film "A.I." presents moviegoers with the latest fictional personification of thinking machines), Cycorp will release to the public a sizable portion of Cyc's so-called knowledge base: the body of assertions and inferences that corresponds to its heart and soul.

The release, under the name OpenCyc, will allow people to use the limited knowledge base for free, browsing it via a Web site or incorporating it in applications ranging from speech-recognition software to database searches.

Users can even supplement the existing knowledge base with new facts and concepts, although these will be subject to review by a technical board. The full knowledge base (about 20 times the size of the public portion) will be licensed to commercial users for a fee. The company also will release two stand-alone applications, one to identify security holes in large computer networks and another designed to answer queries posed in natural language.

Cyc already exhibits a level of shrewdness well beyond that of, say, your average computer running Windows. In one recent demonstration for a Defense Department project, a Cycorp engineer informed the system they would be discussing anthrax. Cyc responded: "Do you mean Anthrax (the heavy metal band), anthrax (the bacterium), or anthrax (the disease)?" Asked to comment on the bacterium's toxicity to people, it replied: "I assume you mean people (homo sapiens). The following would not make sense: People Magazine."

A Lot of Time, a Lot of Money

Getting to that point has not been easy. The project already has consumed an estimated 500 person-years and \$50 million in investments from, among others, the Defense Department, the pharmaceuticals company GlaxoSmithKline, and Microsoft co-founder Paul Allen.

The scale of the project elicits awe-struck appreciation from supporters and critics alike.

"Having an encyclopedic knowledge base that would cover all of common sense is an absolutely critical goal in AI," says Benjamin J. Kuipers, chairman of the University of Texas computer science department. A former advisor to the Cyc project, Kuipers disagrees with some aspects of Lenat's technical approach, but acknowledges: "Doug was the guy with the guts to do this, and he deserves a lot of credit."

The system today encompasses more than 1.4 million assertions--hundreds of thousands of root words, names, descriptions, abstract concepts, and a method of making inferences that allows the system to understand that, for example, a piece of wood can be smashed into smaller pieces of wood, but a table can't be smashed into a pile of smaller tables. As a software program, Cyc is not embodied in any physical thing: A visitor to Cycorp would see only cubicles filled with programmers contemplating conventional computer monitors displaying Cyc's "knowledge."

An intelligent system on this scale, Cycorp officials contend, has countless potentially profitable applications. "My vision of this company is to be the Intel of intelligent software," says Dwight Lodge, chief executive of Cycorp. "I'd like to have [Cyc inside] a whole range of existing applications."

As an experimental project, Cyc emerged as a response to an intellectual crisis in the field of artificial intelligence. Having emerged as a formal science in the mid-1950s, after 30 years, its goals had proved elusive.

Projects from chess-playing computers to robots that could haltingly negotiate uneven terrain took years longer to achieve than expected. Some highly touted AI programs could exceed human performance only in certain narrow fields--such as diagnosing diseases where the choice among possible causes was limited. Lenat, then a member of the Stanford University computer faculty, thought he understood the problem. Expert systems were "not savants, but idiot savants," he said. They lacked the basic information possessed by an average 8-year-old human: Green is not blue; sweating makes you wet; a car can rust but not run a fever. They broke down when they were asked questions that relied, however subtly, on such mundane observations.

The machines had been stuffed with the wrong kind of knowledge. Computers had become repositories of the esoteric facts one found in reference works, textbooks and the brains of experts. What was missing was the vast fabric of ordinary facts and observations that humans acquire and use almost subconsciously and without which life would be unintelligible--what Lenat terms "common sense."

This common sense deficit, Lenat argued, is what makes computer intelligence so shallow. It is why the same system capable of mapping the trajectory of a hurtling ICBM can be brought to a grinding halt by a trivial misspelling in a typed command--one that a child would disregard--or why a computer familiar with human diseases and asked what ails a rusting car is likely to answer: "measles."

Lenat's solution was to program the computer not with the basic information one

finds in the library, but with all the information "the author of an article assumes the reader already knows."

Many people in AI had accepted that the inability to represent common sense was an obstacle, but few had confronted the horrifying necessity of entering so much of it by hand. Lenat resolved to shoulder a task he later called "a 20-year detour to pull the mattress off the road so the traffic can flow."

Whether the result adds up to genuine intelligence, much less the elusive quality we call "consciousness" or "mind," touches on one of the central debates in artificial intelligence. Within the contentious community of AI researchers, Cyc has drawn its share of criticism.

Some critics argue that Cyc's focus on sorting facts and observations into logical categories, even given its powerful ability to make inferences, is too restrictive.

"I don't believe in the idea that intelligence is founded upon having vast amounts of facts about the world," says Douglas R. Hofstadter, the Pulitzer Prize-winning author of "Godel, Escher, Bach: An Eternal Golden Braid," a study of human creativity and artificial intelligence. "Intelligence is about making decisions based on imperfect knowledge and among partially good choices."

But others believe that the exponential surge in the processing power and speed of computers in recent years may finally be enough to give systems such as Cyc the critical mass they need to cross the consciousness threshold.

"We finally are getting to the point where machines will be able to do what the human brain alone can do," says James C. Spohrer, chief technical officer of IBM's venture capital relations group, who has studied Cyc's potential as a commercial project. "The time feels right."

Lenat believes that systems such as Cyc could replicate human cognition closely enough to simulate consciousness, emotion and motivation. Some of these qualities already have appeared in the system, he says.

"Cyc has goals, long- and short-range," he says. "It has an awareness of itself. It doesn't 'care' about things in the same sense that we do, but on the other hand, we ourselves are only a small number of generations away from creatures who operated solely on instinct."

The potential applications of such a discerning system are vast, he says: Systems could converse with their users in plain English or perform accurate translations. Or automated systems could be entrusted with life-and-death responsibilities. "Will you let a robot around your house if it doesn't understand the [relative] value of things, like a moth versus a baby?" Lenat asks.

Cyc already has displayed the ability to identify common-sense absurdities. "Cyc already knows that people have to be a certain age before they're hired for a job," Lenat says, meaning that it could clear such inaccurate entries as mistaken birth dates from corporate payroll records. Cyc also can extract and compile facts scattered among diverse sources of information and use them to draw conclusions--in one test responding to a request for an image of people relaxing by turning up a photo of some men holding surfboards.

An Ongoing Dialogue With Colleague

The center of all this activity is an exceedingly unusual place, even for an emerging technology company. Cycorp's 65-member staff engages in a dialogue day and night with their unremittingly curious electronic colleague.

Most of Lenat's programmers are trained not in computer engineering but in fields related to logic and human thought: The staff includes about 20 philosophers and smaller teams of experts in subjects ranging from theology to physics.

Among them is Charles Klein, 33, a University of Virginia-trained metaphysician who joined Cycorp in 1999 after finding its want-ad for "ontological engineers" in a meager professional quarterly called Jobs for Philosophers.

In a room he shares with a large monitor displaying Cyc's characteristic rows of logical queries and responses, Klein spends hours inculcating the system with such abstract concepts as "belief"--a difficult notion for a computer program to grasp, possibly because it has more to do with point of view than with anything true or

false about the real world.

"People who do this enjoy the process of decoding thought," he says of his daily routine of typing assertions into Cyc's database and replying to the computer's minute requests for clarifications. It is the kind of work that only a specialist could love. "Take the phrase, 'I like to go shopping,'" Klein says. "Connecting each word to a concept is fascinating to any philosopher who's interested in the structure of thought and inference."

One thing everybody agrees on is that Cyc would never have got off the ground, much less kept aloft for 17 years, if not for Doug Lenat.

A former wunderkind of computer science, Lenat is now 50. Brash, barrel-chested and with an unruly mop of black hair, he has a distinctive way of interrupting his technical explanations with a wide smile, as though delighted by his own perspicacity.

"If I have an idea, he's one of the five people I can expect to understand it and see what's wrong with it," says Marvin Minsky, a professor at the Massachusetts Institute of Technology and one of the pioneers of the field.

Lenat burst upon the AI scene in the 1970s with a string of now-legendary programming feats. The first, which became the basis of his Stanford doctorate, was a program called Automated Mathematician, or AM. The program was designed to learn not by being fed a diet of new facts but by "discovery"-- given a handful of starting principles, it was to search for new ones.

AM started with 78 basic concepts such as mathematical sets and 243 "rules of thumb" for making hypotheses, judging the intellectual value of its discoveries on a scale of 0 to 1,000, and so on. If it found something intriguing, such as multiplication, it looked for its inverse, thus discovering division.

Launched into action, the system rapidly evolved into a sophisticated scholar engaged in abstruse speculations in mathematical theory. Every so often AM would latch onto a concept so obscure that Lenat would believe it to be original, only to find that it already had been discovered--often by some real-life theoretical genius.

Eventually, AM acquired an uncommon ailment for a computing system: intellectual exhaustion. Having explored the esoteric reaches of mathematics, AM suddenly downshifted into a preoccupation with rudimentary arithmetic. Finally, with the remark, "Warning! No task on the agenda has priority over 200," the system virtually expired, as though from boredom.

With his next program, Eurisko, Lenat attempted to repair the glitches that had caused AM's ennui. Eurisko's spectacular debut came at the 1981 Traveller Trillion Credit Squadron tournament, a futuristic war game that attracted players nationwide. Having been fed the game's 100-page rule book, Eurisko exploited previously unnoticed loopholes, crafting innovative spaceship designs and ingeniously novel strategies to deploy a small, nimble fleet.

Eurisko easily blew the competition out of the heavens, a feat duplicated the following year. Before the 1983 competition, the organizers informed Lenat that if he chose to enter again, they would cancel the games. He retired gracefully, holding the rank of intergalactic admiral.

In 1984, after he had formulated the design principles that would underlie Cyc, he was invited to set up the project at the Microelectronics & Computer Technology Corp., a research consortium founded in Austin by a group of high-tech companies. (MCC spun off Cycorp in 1995.)

His first step was to design a framework of concepts and categories akin to those of a standard thesaurus. The concept "space," for example, might encompass terrain, elbow room and nothingness, as well as city, country and town. Within each frame his team would place the appropriate common-sense assertions: Terrains can vary, cities have borders, and so on.

As implemented by teams of programmers that occasionally exceeded 100, the knowledge base grew exponentially. Each new assertion had to be categorized and its contradictions with other assertions resolved: Cyc could know both that Dracula was a vampire and that vampires do not exist only by understanding that the first

was true in a fictional world and the second in the physical world. This was done by positioning each assertion in its own "microtheory," or context, such as fiction versus fact.

Although not exactly a secret, Cyc remained an enigma to outsiders. Press reports tended to describe it in terms of the large, glowering electronic brain of popular fancy, rather than a sophisticated software system. A local film crew once published a photo purportedly showing the master computer on the Cycorp premises, a massive hardware unit with blinking lights they had encountered in a storeroom; in fact, it was the building's air conditioner.

Does It Know People Can Run?

Computer professionals given a chance to interact with the growing system often came away excited by its capabilities and disillusioned by its shortcomings. When Vaughan Pratt, a Stanford computer scientist, was allowed a 2 1/2-hour session supervised by a Cyc engineer in 1994 he came away impressed by its ability to flag inconsistencies between databases and perform other selected tasks. But when he attempted to ply Cyc with a number of common-sense questions such as "Can people run?" and "How big is the Earth?" he recalled the system seemed to get tied up searching for relevant facts from within its vast database.

"Cyc didn't know what it knew," Pratt said in a recent interview. "The stumbling block was that there was no mechanism for finding specific facts." (Cycorp programmers subsequently told him they had fixed the flaw.)

But in other tests, Cyc blew away the competition as decisively as Eurisko's space cruisers. In July 1998, the Pentagon put Cyc and a dozen other AI systems through their analytical paces, giving each team a package of 300 pages of abstruse data to program in their systems and following up with a series of complicated strategic queries. Cyc scored better than all the other systems put together, according to the company, leading the Pentagon to make it the core of a new experimental program aimed at developing large knowledge bases.

Now, three years later, Lenat believes Cyc is much closer to fulfilling the role of an intelligent system that augments human capabilities, which after all is the central goal of AI research. "Once you have a truly massive amount of information integrated as knowledge, then the human-software system will be superhuman," Lenat says, "in the same sense that mankind with writing is superhuman compared to mankind before writing."

But confident as he is that Cyc is about to emerge as a truly intelligent machine, Lenat is thinking hard about the responsibilities programmers have to ensure the software works exclusively to humans' advantage.

"HAL killed the ['2001'] crew because it had been told not to lie to them, but also to lie to them about the mission," he observes. "No one ever told HAL that killing is worse than lying. But we've told Cyc."